

SYSTEM FOR HOST-TO-HOST CONNECTIVITY USING FICON PROTOCOL OVER A STORAGE AREA NETWORK

BACKGROUND OF THE INVENTION

5

Technical Field

This invention relates to network connectivity between host computer systems over storage area networks (SANs) using the SAN's network protocol such as the FICON protocol.

10

Description of The Prior Art

Network architecture is a blueprint that is used to guide the design and implementation of networks. These blueprints are used because networks are not fixed at a single point in time, but evolve to accommodate changes in both the underlying technologies upon which they are based as well as changes in the demands placed on them. The primary function of a network is to provide general cost-effective, robust and high-performance connectivity among a large number of computers.

Channels and networks are the two primary ways that data is transferred between devices such as processors, and peripherals such as printers. Channels transfer data through switched or direct point-to-point connections and they work by creating a fixed connection between the source and destination devices until the transfer is complete. Channels transfer data at high speeds and are very economical. Networks, on the other hand are collections of nodes such as processors, print devices and workstations. Connections on networks are slower than those made via channels. Channels work best among few devices, and connect via predefined addresses. Networks, on the other hand, can handle multiple requests among multiple connections. Fibre Channel (FC) is a hybrid of both network and channel methods and combines the best aspects of both.

Host to host connectivity within a high-speed network environment is typically accomplished by using several different types of network media. Network media include Token Ring and Ethernet and other topologies. Token Ring is a physical network technology in which hosts are connected in a ring. A token (bit pattern) circulates around the ring and a given node must possess the token before it is allowed to transmit. Token

Ring has a bandwidth of 4 and 16Mbps (megabits per second). Fiber Distributed Data Interface (FDDI) is a high-speed token ring networking technology designed to run over optical fiber. The cables used for the token ring network are attached to adapters instead of the hosts themselves. Ethernet is a network technology that uses CSMA/CD (Common Sense Multiple Access) and has a bandwidth of 10Mbps or 100Mbps. The Ethernet protocol is implemented in hardware on the network adapter.

A SAN is a dedicated, centrally managed, secure information infrastructure which enables any-to-any interconnection of servers and storage systems. SANs allows for connection between storage devices and host machines across greater distances than are possible on a traditional local area network (LAN). SANs enable users to store large volumes of data at remote locations called libraries. SANs are rapidly being integrated into distributed environments using Fibre Channel technology.

Fibre Channel (FC) is a high speed (1.06 Gbps) medium used for data transfer and storage. It is essentially a serial data channel created over fiber optic cabling. Fibre Channel provides a logical bi-directional, point to point connection between a host and a device. Just as existing networking technologies using LAN/WAN (Local-Area-Network / Wide-Area-Network) utilize configurations connected to PCs, servers, print and storage devices, so too does FC. However, because FC uses fiber optic cabling, connections along a FC network are possible. Another feature of FC is its ability to provide increased bandwidth over communications channels.

ESCON (Enterprise Systems Connection) is a fiber-based protocol. The ESCON protocol allows various emulations to run on top of it for further communication with the mainframe either directly or indirectly through ESCON emulators. However, in order to perform this connectivity, the non-mainframe hosts need a separate ESCON hardware adapter, as well as other limitations.

ESCON based physical channel speeds are 17 Mbps. ESCON network connectivity is usually accomplished through ESCON directors and hence the distances between the end-hosts are limited. ESCON requires a repeater every 1.86 miles. This requirement leads to increased costs when the machines or hosts that are to be connected are far apart.

By comparison, FICON (Fibre Connectivity) Channel is a protocol that is targeted to provide high-speed storage Input/Output (I/O) capability. FICON provides a theoretical

throughput of 100Mbytes/sec, which allows it to be used in a SAN (Storage Area Network) environment.

Addressing the need for a centralized storage management solution across heterogeneous networked resources, FC SAN architecture benefits a typical networked environment by delivering a dedicated storage environment, and by using of high-speed fiber-optic or copper cabling to interconnect servers, workstations and storage devices.

As a data-channel network standard, the FC backbone of a SAN environment contains the network features that provide connectivity, distance and protocol multiplexing.

Because FC SANs are protocol-independent, the support for network protocols like Transmission Control Protocol / Internet Protocol (TCP/IP) and audio/video is accomplished with a single physical layer. FC allows concurrent communication among workstations, mainframes, servers, data storage systems and other peripherals using SCSI and IP protocols. The interconnectivity of computers and storage devices forms the basis of a SAN.

Although FC SAN is a fiber-based solution for high-speed storage access, there remains an industry need for high-speed, host to host connectivity that is an improvement upon the existing ESCON based architecture.

SUMMARY OF THE INVENTION

The invention comprises a system and method for providing host-to-host connectivity through use of SAN. In the preferred embodiment, the end-to-end host connectivity is accomplished through the FICON protocol, over the Fibre Channel layered stack on a standard Fibre Channel adapter that runs transparently across the SAN infrastructure.

Other features and advantages of this invention will become apparent from the following detailed description of the presently preferred embodiment of the invention, taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram of a SAN connection according to the prior art.

FIG. 2 is a diagram of SNA & TCP/IP over SAN.

FIGS. 3 and 4 are diagrams of host-to-host connectivity according to the present invention.

FIGS. 5-7 are block diagrams of prior art architecture.

FIGS. 8-9 are block diagrams of the architecture of the present invention.

FIG. 10 is a schematic flow diagram of 3172 emulation according to the present invention.

5 FIG. 11 is a CD-ROM bearing a software embodiment of the invention.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Overview

10 Host-to-host connectivity through SAN offers several technological advantages over the prior art. The new architecture is capable of higher bandwidths and, therefore, higher speeds. There is no longer a need for new FICON hardware, but rather the new architecture uses industry-proven FC hardware. Hardware proofing cycles are eliminated, and there is no device driver development required in the present invention. Therefore, resulting in a
15 significant reduction in the development life-cycle of the product. The new architecture uses the existing emulation software layers and Fibre Channel stack, and the new architecture offers improved connectivity at reduced costs.

Technical Background

20 FIG. 1 depicts the prior art connectivity that utilizes a SAN. Host A 1 is able to share data storage 6 with Host B 2, which could be a mainframe through the SAN 5. Host A 1 utilizes the SCSI protocol over FC 3 while communicating with SAN 5, while Host B 2 utilizes the FICON protocol over FC 4. However, this configuration does not support direct Host A to Host B communications.

25 The prior art network architecture that would support the connection 3 between Host A 1 and the SAN 5 of FIG. 1 is depicted in FIG. 6. The physical fiber 39 are the cables preferably fiber optic. The fibre channel stack 38 is a protocol layered stack and SCSI 37 is a protocol layer on top of the fibre channel stack. The Physical fiber and stacks provide the means or channel for communications to the storage access 36.

30 The prior network configuration that would support direct communications between two hosts is depicted in FIG. 2. In order for Host A 8 to communicate directly with Host B 9 in either a local area network (LAN environment or wide-area-network (WAN) and still maintain access with storage 12 through the SAN 11. The SNA & TCP/IP protocol channel 10. Host A 8 and Host B 9 are connected to storage 12 through the SAN 11.

35 FIG. 7 depicts the architecture for both the SAN and LAN/WAN networks. The architecture for the SAN utilizes the FICON protocol 41. The LAN or WAN network

connection between Host A 8 and Host B 9 is facilitated by the architecture consisting of the Ethernet or Token Ring 47, TCP/IP 46, SNA 45 and the applications 44. Note the additional hardware and protocols that are needed in order to provide host to host communication and to access applications.

FIG. 3 depicts the network configuration according to the present invention that facilitate direct communication between two hosts through a SAN. The SAN 17 utilizes an architecture consisting of TCP/IP and system network architecture (SNA) over FICON. Host A 14 utilizes SCSI & FICON protocols over FC. Host B utilizes FICON protocol over FC. This connection scheme allows direct communication between Host A 14 and Host B 15 through the SAN 17.

FIG. 4 is an example connection scheme that utilizes the present invention. The NT 21, Unix 22 and MVS 23 systems are used in this example of the preferred embodiment. Assume that the NT 21, UNIX 22 and MVS 23 systems are connected to the SAN 25 fabric through a standard Fiber Channel connection. On the UNIX system side, a separate FC adapter is dedicated for the host connectivity in order to avoid the FC cable / port sharing with the storage I/O activity on this system. The FICON protocol 24 runs over the Fibre Channel stack, and it is further required that it runs on both of these system objects of interest. After the FC2 level connection is established, the FICON protocol then runs end-to-end between these two systems, thus offering its services to the emulation layers on top of it. These emulation layers further provide the services to TCP/IP, SNA etc. which, ultimately service the applications. The layering architecture of the present invention is explained in FIG. 8.

In the preferred embodiment the SAN network is intended not only to perform storage IO, but the SAN can also perform host-to-host network connectivity without interfering with the storage IO activity. In the preferred embodiment the FICON protocol runs transparently as far as the SAN is concerned, and therefore, the SAN is not aware of the type of traffic that has been passing through. As shown in FIG. 3, the primary purpose of the FICON protocol is to perform the storage IO between the MVS systems and their storage systems, either directly or indirectly, through the SAN infrastructure at a much higher bandwidth than it is possible with ESCON.

The network architecture of the present invention as shown in FIG. 8, provides a layered stack concept logically extending the FC layered stack. The new network architecture is a combination of the architectures in FIGs. 5 and 7. The physical fiber 53 embodies the network hardware. The fiber channel stack 52 is above the physical fiber. The FICON protocol 51 facilitates host communication with the SAN. The emulation layers 50

facilitate the imitation of other programs or devices. The protocol stacks 49 facilitate host-to-host communication and access to applications 48.

FIG. 9 is an example of the network architecture of the present invention. The FICON protocol 61 must run on top of the FC stack 62 which allows it to support different types of emulation such as, 3172 57, 3088 59 or 3174 60. Each emulation supports its relative function, such as TCP/IP 56 or high-speed file transfer application 55. After establishing a FICON level end-to-end connection, the respective emulations on both sides can exchange the communication for further support of their service users.

FIG. 10 is an example of how an emulation layer would interface with the layer above and below it and in this example, 3172 emulation is considered for a detailed explanation of how it works. The 3172 emulation layer 57 utilizes the FICON services at the lower interface for communicating with its peer entities. However, the communication with the remote peer entities is performed through IBM standard 8232 protocol to support TCP/IP pass-through. At the upper interface of this emulation layer, a set of service primitives are provided so that the user of this emulation layer such as the TCP/IP stack in this particular case, could communicate using these primitives. The emulation layer 57 encapsulates the TCP/IP packets in 8232 protocol frames 66 with the 8232 protocol header 67 attached before transmitting them to a remote peer entity or services. The 8232 protocol frames 67 are then transmitted using the FICON layer. Similarly, the emulation layer will strip the 8232 header 68 from the received 8232 protocol frames from the underlying FICON layer before handing over the encapsulation and decapsulation, the emulation layer also handles the control frames to manage the 8232 protocol activity.

The 3172 emulation could support Ethernet, FDDI and Token Ring LAN types and accordingly, it exposes the appropriate interfaces to the TCP/IP stack above it.

Advantages Over The Prior Art

The present invention has an advantage over the prior art in that no special FICON based hardware adapters are required, unlike ESCON based architecture, in order to provide end-to-host network connectivity. Instead the FICON based architecture uses standard Fibre Channel adapters. This allows the SAN infrastructure to be leveraged for physical connectivity needs. Also, the distance between the end-hosts is no longer a limiting factor from a cost perspective because the connectivity can be easily extended at a very reduced price through WAN networks. Since no special hardware other than the standard Fibre Channel hardware adapter is required, there is a significant reduction in product development or test cycles. Also there is no additional cost associated with software driver

development for the adapter. Ultimately the present invention will result in better time-to-market (TTM) of products, significant cost reductions and greater product flexibility.

Alternative Embodiments

5 FIG. 11 shows an article of manufacture or a computer program product including a signal bearing media for storing thereon program means for carrying out the method of this invention in the system of FIGS. 3 and 8. While the present invention has been described in the context of a system and method of operation, the mechanisms of the present invention are capable of being distributed as a program-product in a variety of forms, and the present
10 invention applies equally regardless of the particular type of signal bearing media used to actually carry out the distribution. Examples of signal-bearing media include recordable media (such as digital and analog communications links, including fiber and wireless communications links).

15 An example of such an article of manufacture is illustrated in FIG. 11 as pre-recorded CD-ROM 1002. CD-ROM 1002 is intended for use with a data processing system, and includes optical storage medium 1004, and program means 1006, 1008, 1010 and 1012 recorded thereon for causing the system of FIGS. 3 and 8, to practice steps of this invention. It will be understood that such apparatus and articles of manufacture all fall within the spirit and scope of this invention.

20 It will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without departing from the spirit and scope of the invention. In particular, other protocols could be used over the SAN / FICON architecture. Accordingly, the scope of protection of this invention is limited only by the following claims and their equivalents.